

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problems Mailbox.**

THIS PAGE BLANK (USPTO)

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F-17/30		A1	(11) International Publication Number: WO 98/41932
			(43) International Publication Date: 24 September 1998 (24.09.98)
(21) International Application Number: PCT/FI98/00191			(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 4 March 1998 (04.03.98)			
(30) Priority Data: 971066 14 March 1997 (14.03.97) FI			
(71) Applicant (for all designated States except US): NOKIA TELECOMMUNICATIONS OY [FI/FI]; Keilalahdentie 4, FIN-02150 Espoo (FI).			
(72) Inventors; and (75) Inventors/Applicants (for US only): TIKKANEN, Matti [FI/FI]; Viherlaaksonranta 10 A 1, FIN-02710 Espoo (FI). IIVONEN, Jukka-Pekka [FI/FI]; Wallininkatu 6 C 72, FIN-00530 Helsinki (FI).			
(74) Agent: PATENT AGENCY COMPATENT LTD.; Teollisuuskatu 33, P.O. Box 156, FIN-00511 Helsinki (FI).			

Published

With international search report.


Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

In English translation (filed in Finnish).

(54) Title: METHOD FOR IMPLEMENTING AN ASSOCIATIVE MEMORY BASED ON A DIGITAL TRIE STRUCTURE

(57) Abstract

The invention relates to a method for implementing a memory. The memory is implemented as a directory structure comprising a tree-shaped hierarchy having nodes at several different levels, wherein an individual node can be (i) a trie node comprising an array wherein an individual element may contain the address of a lower node in the tree-shaped hierarchy and wherein an individual element may also be empty, or (ii) a bucket containing at least one element so that the type of an individual element in the bucket is selected from a group including a data unit, a pointer to a stored data unit, a pointer to another directory structure and another directory structure. To minimize storage space requirement, the arrays in the trie nodes are implemented as quad nodes having a fixed size of four elements, and in at least part of the directory structure sets of successive quad nodes are replaced with compressed nodes in such a way that an individual set made up by successive quad nodes, from each of which there is only one address to a quad node at a lower level, is replaced with a compressed node (CN) storing an address to the quad node that the lowest node in the set to be replaced points to, information on the value of the search word by means of which said address is found, and information on the total number of bits from which search words are formed in the set to be replaced. The invention also relates to a structure in which buckets are not employed.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon	KR	Republic of Korea	PL	Poland		
CN	China	KZ	Kazakhstan	PT	Portugal		
CU	Cuba	LC	Saint Lucia	RO	Romania		
CZ	Czech Republic	LI	Liechtenstein	RU	Russian Federation		
DE	Germany	LK	Sri Lanka	SD	Sudan		
DK	Denmark	LR	Liberia	SE	Sweden		
EE	Estonia			SG	Singapore		

Method for implementing an associative memory based on a digital trie structure

Field of the Invention

The present invention generally relates to implementation of an associative memory, particularly to implementation of an associative memory based on a digital trie structure. The solution in accordance with the invention is intended for use primarily in connection with central memory databases. Suitable applications particularly include databases in which there are a large number of insertions and deletions compared to retrievals. Such applications include visitor location registers (VLR) in mobile communications networks.

Background of the Invention

The prior art unidimensional directory structure termed digital trie (the word "trie" is derived from the English word "retrieval") is the underlying basis of the principle of the present invention. Digital tries can be implemented in two types: bucket tries, and tries having no buckets.

A digital bucket trie structure is a tree-shaped structure composed of two types of nodes: buckets and trie nodes. A bucket is a data structure containing a number of data units or a number of pointers to data units or a number of search key/pointer pairs (the number may include only one data unit, one pointer or one key/pointer pair). A trie node, on the other hand, is an array guiding the retrieval, having a size of two by the power of k (2^k) elements. If an element in a trie node is in use, it refers either to a trie node at the next level in the directory tree or to a bucket. In other cases, the element is free (empty).

Search in the database proceeds by examining the search key (which in the case of a subscriber database in a mobile telephone network or a telephone exchange, for instance, is typically the binary numeral corresponding to the telephone number of the subscriber) k bits at a time. The bits to be searched are selected in such a way that at the root level of the structure (in the first trie node), k leftmost bits are searched; at the second level of the structure, k bits next to the leftmost bits are searched, etc. The bits to be searched are interpreted as an unsigned binary integer that is employed directly to index the element array contained in the trie node, the index indicating a given element in the array. If the element indicated by the index is free, the search will terminate as unsuccessful. If the element refers to a trie node at the

next level, k next bits extracted from the search key are searched at that level in the manner described above. As a result of comparison, the routine branches off in the trie node either to a trie node at the next level or to a bucket. If the element refers to a bucket containing a key, the key stored therein is compared with the search key. The entire search key is thus compared only after the search has encountered a bucket. Where the keys are equal, the search is successful, and the desired data unit is obtained at the storage address indicated by the pointer of the bucket. Where the keys differ, the search terminates as unsuccessful.

10 A bucketless trie structure has no buckets, but reference to a data unit is effected from a trie node at the lowest level of a tree-shaped hierarchy, called a leaf node. Unlike buckets, the leaf nodes in a bucketless structure cannot contain data units but only pointers to data units. Also a bucket structure has leaf nodes, and hence trie nodes containing at least one pointer to a bucket (bucket structure) or to a data unit (bucketless structure) are leaf nodes. The other nodes in the trie are internal nodes. Trie nodes may thus be either internal nodes or leaf nodes. By means of buckets, the need for reorganizing the directory structure can be postponed, as a large number of pointers/data units can be accommodated in the buckets until a time when the need for reorganization arises.

20 The solution in accordance with the invention can be applied to a bucket structure as well as a bucketless structure. In the following, bucket structures will nevertheless be used as examples.

Figure 1 illustrates an example of a digital trie structure in which the key has a length of 4 bits and $k=2$, and thus each trie node has $2^2=4$ elements, and two bits extracted from the key are searched at each level. Buckets containing a record are denoted with references A, B, C, D...H...M, N, O and P. Thus a bucket is a node that does not point to a lower level in the tree. Trie nodes are denoted with references IN1...IN5 and elements in the trie node with reference NE in Figure 1.

30 In the exemplary case of Figure 1, the search keys for the buckets shown are as follows: A=0000, B=0001, C=0010,..., H=0111,... and P=1111. In this case, a pointer is stored in each bucket to that storage location in the database SD at which the actual data, e.g. the telephone number of the pertinent subscriber and other information relating to that subscriber, is to be found. The actual subscriber data may be stored in the database for instance as a

sequential file of the type shown in the figure. The search is performed on the basis of the search key of record H, for example, by first extracting from the search key the two leftmost bits (01) and interpreting them, which delivers the second element of node IN1, containing a pointer to node IN3 at the next level.

- 5 At this level, the two next bits (11) are extracted from the search key, thus yielding the fourth element of that node, pointing to record H.

Instead of a pointer, a bucket may contain (besides a search key) an actual data file (also called by the more generic term data unit). Thus for example the data relating to subscriber A (Figure 1) may be located in bucket A, the data relating to subscriber B in bucket B, etc. Thus in the first embodiment of an associative memory, a key-pointer pair is stored in the bucket, and in the second embodiment a key and actual data are stored, even though the key is not indispensable.

10 The search key may also be multidimensional. In other words, it may comprise a number of attributes (for example the family name and one or more forenames of a subscriber). Such a multidimensional trie structure is disclosed in international application No. PCT/FI95/00319 (published under number WO 95/34155). In said structure, address computation is performed in such a way that a given predetermined number of bits at a time is selected from each dimension independently of the other dimensions. Hence, a fixed limit independent of the other dimensions is set for each dimension in any individual trie node of the trie structure, by predetermining the number of search key bits to be searched in each dimension. With such a structure, the memory circuit requirement can be curbed when the distribution of the values of the search keys is known in advance, in which case the structure can be implemented in a static form.

25 It has been found in practice, however, that an optimum fill rate for trie nodes in view of storage space occupancy cannot always be achieved with this structure (with all key distributions).

30

Summary of the Invention

It is an objective of the present invention to remove the above drawback and to provide a method whereby, in a structure employing fixed node size, memory space requirement can be maintained as small as possible and memory management can be maintained effective, particularly in databases in which there are a large number of insertions and deletions.

35

This objective is achieved with the solution defined in the independent claims. The first of these discloses a structure employing buckets and the second a structure not employing buckets.

The idea of the invention is to use in the structure a fixed node size of four (elements) and to replace successive quad nodes providing only one path downward in the tree with (one or more) compressed node(s) in which the information needed for proceeding in the structure and for restructuring the nodes is collected from one or more successive quad nodes, without any storage space being needed for element arrays. Thus, the idea is to use an optimum node size in view of storage space requirement and memory efficiency, and moreover to further diminish storage space requirement for successive quad nodes providing only one path downward in the tree (and not having an optimum fill rate). In other words, the idea is to curb storage space requirement particularly in conjunction with non-uniform key distributions by reducing the depth of the structure by means of compressed nodes.

The solution in accordance with the invention also ensures effective performance of set operations, as the structure is an order-preserving digital trie.

Brief Description of the Drawings

In the following the invention and its preferred embodiments will be described in closer detail with reference to examples in accordance with the accompanying drawings, in which

- Figure 1 illustrates the use of a unidimensional digital trie structure in the maintenance of subscriber data in a telephone exchange,
- Figure 2 shows a multidimensional trie structure,
- Figure 3 shows a three-dimensional memory structure,
- Figure 4 illustrates a memory structure in accordance with the invention,
- Figure 5 illustrates implementation of address computation in the memory of the invention,
- Figures 6a and 6b illustrate replacement of quad nodes with compressed nodes,
- Figures 7a...7c show the principle of forming a compressed node,
- Figures 8a and 8b show an example of the maintenance of the memory structure,

Figure 9 illustrates the structure of a compressed node employed in the memory,

Figure 10a illustrates the limitation posed by the word length employed on combining the nodes,

5 Figure 10b shows the structure of a collecting node to be formed from the node chain of Figure 10a, and

Figure 11 shows the memory arrangement in accordance with the invention on block diagram level.

10

Detailed Description of the Invention

As stated previously, in the present invention the trie structure has a multidimensional (generally n-dimensional) implementation. Such a multidimensional structure is otherwise fully similar to the unidimensional structure described at the beginning, but the element array contained in the trie node is multidimensional. Figure 2 exemplifies a two-dimensional $2^2 \times 2^1$ structure, in which one dimension in the element array comprises four elements and the other dimension two elements. Buckets pointed to from the elements in the trie node are indicated with circles in the figure.

Address computation in the multidimensional case is performed on the same principle as in the unidimensional case. The fundamental difference, however, resides in that instead of a single element array index, an index is calculated for each dimension in the element array (n indices). Each dimension thus has a search key space of its own $\{0, 1, \dots, 2^{v_i}-1\}$ (v_i is the length of the search key in bits in each dimension and $i \in \{1, \dots, n\}$).

25 The size of the trie node in the direction of each dimension is 2^{k_i} elements, and the total number of elements S in the trie node is also a power of two:

$$S = \prod 2^{k_i} = 2^{k_1} \times 2^{k_2} \times 2^{k_3} \times \dots = 2^N \quad (1)$$

30 All elements in a trie node having n dimensions can thus be pointed to by n integers ($n \geq 2$), each of which may have a value in the range $\{0, 1, \dots, 2^{k_i}-1\}$. Thus the predetermined fixed parameter is the total length of the search key in each dimension. If for example one dimension of the search key has 256 attributes (such as first names) at most, the total length of the search key is 8 bits.

35 Figure 3 shows an example of a node N10 used in the directory structure of the memory, employing a three-dimensional search key. In the

direction of the first dimension (x), the trie node has $2^2=4$ elements, in the direction of the second dimension (y) $2^1=2$ elements, and in the direction of the third dimension (z) $2^3=8$ elements, which gives a total of $2^6=64$ elements in the trie node, numbered 0...63.

5 Since the memory space in practical hardware implementations (for example computer equipment) is unidimensional, the multidimensional array is linearized, i.e. converted to be unidirectional, in the address computation operation (that is, in proceeding in the directory tree). The linearization is an arithmetic operation that can be performed on arrays of all sizes. Hence, it is irrelevant whether the trie nodes are considered to be unidimensional or multi-
10 dimensional, as multidimensional arrays are linearized in any case to be unidimensional.

In linearization, the elements in the array are numbered starting from zero (as shown in Figure 3), the number of the last element being on less than the product of the sizes of all dimensions. The number of an element
15 is the sum of the products of each coordinate (for example in the three-dimensional case, the x, y and z coordinates) and the sizes of the dimensions preceding it. The number thus computed is employed directly to index the unidimensional array.

20 In the case of the trie node shown in Figure 3, the element number VA_n is calculated in accordance with the above with the formula:

$$VA_n = x + y \times 4 + z \times 4 \times 2 \quad (2)$$

where $x \in \{0,1,2,3\}$, $y \in \{0,1\}$ and $z \in \{0,1,2,3,4,5,6,7\}$. Thus for example for element 54 we obtain from the coordinates thereof (2,1,6):
25 $2 + 1 \times 4 + 6 \times 4 \times 2 = 2 + 4 + 48 = 54$.

When the (n-dimensioned) element array of a trie node of an n-dimensional trie structure is linearized, in accordance with the above the size of each dimension is 2^{k_i} , where k_i is the number of bits to be searched at a time in the dimension concerned. If a coordinate in accordance with the dimension
30 is denoted by reference a_j ($j \in \{0,1,2,\dots,n\}$), the linearization can be written out as

$$\sum_{j=1}^n a_j \prod_{i=0}^{j-1} 2^{k_i}, \forall j: a_j \in \{0,1,2,\dots,2^{k_j}-1\} \wedge k_0 = 0 \quad (3)$$

The linearization can be carried out by performing a multiplication in accordance with formula (3); yet it is expedient to perform the linearization by forming from the search key bits a bit string by known methods, the corre-

sponding numeral indicating the element whose content provides the basis for proceeding in the directory tree. Such a linearization method is termed bit interleaving.

As stated previously, it has been found in practice that the solution disclosed in the international application referred to at the beginning is not capable of achieving optimum storage space occupancy in all conditions (with all key distributions). When this fact is studied mathematically, it can be shown that the average storage space occupancy E of the trie nodes in an n -branched trie structure (n -ary trie) with uniform key distributions is in accordance with Formula (4)

$$E = \frac{n \log e}{m \log n} v - \frac{n}{n-1} \quad (4)$$

where n is the size of the array in the trie node, m is the maximum number of pointers in the bucket, v is the number of records in the entire structure, and e is neper. Formula (4) can be deduced from the formulae disclosed in the article C. H. Ang and H. Samet, *Approximate average storage utilization of bucket methods with arbitrary fanout*, Nordic Journal of Computing 3 (1996), pp. 280-291 (also *Computer Science TR 3618*).

In Formula (4), the essential factor for the invention is that it has minima for values 2 and 4 of parameter n . Thus, storage space occupancy is at its minimum with values 2 and 4. Of these values, value 4 is selected for use in the memory in accordance with the invention, since this will allow smaller depth in the structure and higher efficiency (rapidity) in the memory. In the present invention, therefore, the trie structure is implemented as a dimension-independent 4-branched (quad) structure. This means that the number n of the above-described dimensions may have any value, yet the number of bits to be searched in the trie nodes is fixed, i.e. two (with the exception of compressed nodes, which will be described later). In other words, four is selected as the fixed size of all ordinary trie nodes in the structure, regardless of the number of dimensions employed. Such trie nodes will also be termed quad nodes hereinbelow.

One such structure is illustrated in Figure 4. The figure depicts five ordinary quad nodes $N41...N45$, having altogether six pointers to buckets ($L1...L6$). In addition to quad nodes and buckets, the structure may also have compressed nodes, which will be described in detail later.

Figure 5 illustrates an example of address computation performed in trie nodes in the memory in accordance with the invention. It has been presumed in the figure that the space is three-dimensional (dimensions x, y and z), but as stated above, the number of spatial dimensions can be arbitrary. It has further been presumed that search key a_x in the direction of dimension x is $a_x = 011011$, search key a_y in the direction of dimension y is $a_y = 110100$ and search key a_z in the direction of dimension z is $a_z = 101010$. The search keys are listed in vertical sequence in the figure.

In the nodes of the trie structure, the indexing bits of the unidimensional element array are shown in frames denoted by continuous lines. These frames illustrate how a global search key is divided into local search keys, each being used in one node of the trie structure. The nodes in the structure are denoted by references N1...N9 in the order of progress. Two bits are searched in each node, and thus in the first node the leftmost bit of search key a_x and the leftmost bit of search key a_y are used. Thereafter the routine proceeds in the direction of the arrow to the next node (N2), in which the element array index is formed by the next two bits. These are the leftmost bit in search key a_z and the next to the leftmost bit in search key a_x . Proceeding in this way, the following node-specific element array indices are obtained: 01 (node N1), 11 (node N2), 10 (node N3), 10 (node N4), 10 (node N5), 10 (node N6), 10 (node N7), 11 (node N8) and 00 (node N9).

Since the number of bits to be searched in each trie node is constant, i.e. two, the number of searched bits need not be stored in the node. In its minimum configuration, the node thus only comprises an element array having four elements. For proceeding in the directory tree, the type of each node must additionally be known. This data can be stored in the directory structure for example in each node or in the pointer of the parent of the node. By means of the two "extra" bits of the pointer, information can be encoded in the pointer on whether a zero pointer (an empty element) is concerned or whether the pointer points to an ordinary trie node, a bucket or a compressed trie node (which will be described hereinbelow). In the case of a bucketless structure, information on whether the pointer points to an uncompressed node, a compressed node or a data unit is stored.

Since the address computation is performed by using bit interleaving known per se, the address computation will not be described in further detail.

To further minimize storage space occupancy, compressed nodes are formed from the quad nodes in certain cases. If an ordinary quad node has only one child, this means that only one path "downward" in the tree passes through said quad node. To optimize memory space occupancy, a quad node containing only a single pointer (path downward) should preferably be replaced with a compressed node in which the number of bits searched in said path and the computed array index value are disclosed. Since it is advantageous from the point of view of storage space requirement to form compressed nodes from single-child quad nodes throughout the entire memory structure, compression also means that at least two child nodes are always maintained for ordinary (uncompressed) quad nodes in the memory structure, that is, an individual (ordinary) quad node has pointers to at least two different lower-level nodes (child nodes). A compressed node replaces one or more successive internal nodes, each of which has one child, and hence the above-stated one child cannot be a bucket (or a leaf in a structure that has no buckets). Hence, a child node must be an ordinary quad node in order for compression to be possible. From the point of view of optimizing storage space, it is thus advantageous to always maintain at least two child nodes for quad nodes preceding a bucket as well.

Figures 6a and 6b illustrate the principle of forming compressed quad nodes. If a situation in accordance with Figure 6a in which only one path passes through one or more successive quad nodes QN is created in the structure, the quad nodes in said path can be replaced with one compressed node CN (Figure 6b), having information on the number of searched bits B, an element array index i, and a pointer to a child node. This child node is always an ordinary quad node. In other words, the compressed node does not point to a bucket. If the number of search key bits searched in said path is greater than the word length used in the computer employed, more than one compressed nodes are needed to replace the quad nodes in said path.

The memory in accordance with the invention thus comprises two types of trie nodes: ordinary quad nodes containing an element array in accordance with Figure 5, and compressed nodes that will be described in the following.

Figures 7a and 7b illustrate the forming a compressed quad node. Figure 7a shows a local trie structure made up by three quad nodes, providing only one path for the two uppermost nodes. These two nodes can be replaced

with one element array shown in Figure 7b. Since the structure has a single path for said nodes, only one element of the array is in use, which in this exemplary case is element 14 circled in the figure (14=1011 when the bits are taken in line order, i.e. the x bits first and thereafter the y bits). Thus, for the
5 two uppermost nodes the trie structure of Figure 7a can be replaced with a compressed quad node in which the number of bits to be searched (4) and the value of the array index (14) are stored.

Figures 8a and 8b show a maintenance example corresponding to the previous figure when data units and associated keys are deleted from a
10 database. Figure 8a shows an initial situation in which the memory structure comprises trie nodes N81...N83 and buckets L2...L4. Thereafter bucket L2 and the pointer/record contained therein is deleted from the memory, as a result of which nodes N81 and N82 can be replaced with a compressed node CN, in which the index of the pointer contained in the node and the number of
15 bits searched in the path replaced by the compressed node are disclosed.

Hence, the compressed node is in principle similar to an ordinary quad node, but instead of an element array with only one pointer being stored, the index of the pointer concerned and the number of bits searched in the path are stored. A compressed node thus has a virtual array replacing the information contained in the one or more node arrays existing in the path. If the compressed node replaces several ordinary quad nodes, the number of searched
20 bits indicated in the compressed node is equal to the sum of the numbers of bits searched in the replaced nodes.

Figure 9 illustrates the structure of a compressed node. The minimum configuration of the node comprises 3 parts: field 90 indicating the number of searched bits, field 91 storing the value of the array index, and field 92 storing a pointer to a child node. The compressed node is in need of this data in order for the search to proceed with the correct value at the compressed node as well, and in order for the restructuring of the node to be possible in
30 connection with changes in the structure. (Without information on the number of searched bits, the array index value cannot be calculated from the search key, and on the other hand without the array index value the calculated value could not be compared to the value stored in the node.)

If a collision occurs in the compressed node in connection with an
35 insertion, i.e. the compressed node will have a new pointer, it is studied which bit in order distinguishes the index of the initial pointer and the index of the

new pointer. Accordingly, a structure replacing the initial compressed node is created, in which the new compressed node comprises the number of index bits insofar as there are common bits. In addition, one or more trie nodes are created in the structure at points corresponding to those bits in which the indices differ from one another.

5 If the compressed node is preceded by one or more compressed nodes or a chain of trie nodes providing only a single path, it is advantageous in view of storage space requirement and memory efficiency to further combine said nodes. Moreover, in view of memory efficiency it is advantageous to
10 carry out the combination of nodes in such a way that only in the compressed node that is the last (lowest) in the chain the number of searched bits is smaller than the word length in the computer used. In other words, nodes are combined in such a way that the number of searched bits will be as large as possible in each compressed node. For example, three successive compressed nodes in which the numbers of searched bits are 5, 10 and 15 can be
15 combined into one compressed node in which the number of searched bits is 30. Likewise, for example three successive compressed nodes (or three successive ordinary trie nodes providing only one path) can be combined into two compressed nodes in which the numbers of searched bits are 32 and 3, with
20 the word length employed being 32. Hence, it is attempted to obtain in as many compressed nodes as possible a number of searched bits corresponding to the word length of the computer, and the possible "superfluous" bits are left for the compressed node that is lowest in the hierarchy.

However, compressed nodes cannot be combined so as to make
25 the number of bits searched in one node higher than the word length in the computer employed. Particularly in multidimensional cases ($n > 3$), it is common that there are so many successive nodes containing one child that the path cannot be represented by a single compressed node. Therefore, the search path or part thereof is replaced with a chain made up by several successive
30 compressed nodes, in which the number of searched bits is the same as the word bit number, for example 32 in the Intel architecture, except for the last node where the number of bits is smaller than or equal to the word bit number.

Such a situation is depicted in Figure 10a, showing three successive compressed nodes CN1...CN3. The numbers of bits searched in the
35 nodes are denoted by references b, b' and b" and the values of the array indices contained in the nodes with i, i' and i", respectively. In the two uppermost

nodes, the number of searched bits has a maximum value (providing that a 32-bit computer architecture is used).

It is advantageous to form from a chain of several successive compressed nodes resulting from limited word length a single node collecting such compressed nodes. This collecting node is formed in such a way that the pointer of the collecting node is set to point to the child of the compressed node that is last in said chain, the sum of the numbers of bits searched in the compressed nodes in the chain is set as the number of bits B searched in the collecting node, and the array indices (i.e. search words) produced by bit interleaving are inserted in the list or table T of the node in the order in which they appear in the successive compressed nodes. Thus, the collecting node will be a node CN4 as shown in Figure 10b, comprising three parts: field 130 containing a pointer to said lower-level node, field 131 containing the number of searched bits B (the above sum), and list or table T containing in succession the array indices produced by bit interleaving. This third part thus has a varying size. In the example of the figure, the number of indices is three, since the example of Figure 10a comprises three successive nodes.

The number of elements (i.e., indices) NE in table T is obtained from the number of searched bits B as follows:

$$NE = \begin{cases} B/W, & \text{if } B \bmod W = 0 \\ \lfloor B/W \rfloor + 1, & \text{if } B \bmod W \neq 0 \end{cases}$$

where $\lfloor \cdot \rfloor$ is a floor function truncating decimals from the number, W is the word length used, e.g. 32, and MOD refers to modulo arithmetic. Thus, the number of indices need not be stored in the collecting node as separate data, but it can be found on the basis of the number of searched bits.

The number of bits B' needed to calculate the last index in the table (denoted by reference b'' in the figure), which does not necessarily equal the word length, is obtained as follows:

$$B' = \begin{cases} W, & \text{if } B \bmod W = 0 \\ B \bmod W, & \text{if } B \bmod W \neq 0 \end{cases}$$

By forming a collecting node from several successive compressed nodes, the number of memory references (pointers) can be reduced further. In present-day computer architecture, comprising caches of various levels, memory references require considerable computation time, and hence the compu-

tation time will be diminished. At the same time, the need for storage space for pointers is eliminated.

By means of compressed nodes, the storage requirement can be effectively minimized particularly in conjunction with non-uniform key distributions, since by means of compression the depth of the structure can be arbitrarily increased on a local basis without increased storage space requirement.

As was already indirectly stated above, in the memory in accordance with the invention a bucket cannot be preceded by a compressed node, but the parent node is always either an ordinary quad node or an empty element. Hence, a compressed node cannot point to a bucket, but it always points either to another compressed node or to an ordinary quad node. An empty element means that if the total number of records is smaller than the number of pointers/records that the bucket can accommodate, a tree-shaped structure is not needed yet, but one bucket will suffice in the structure (in which case said node is conceptually preceded by an empty element). It is advantageous to proceed in this way at the initial phase of starting up the memory. It is thus worth-while starting building up the tree-shaped structure only when this is necessary.

In other respects, the retrievals, insertions and deletions performed in the memory are made in a manner known per se. In this regard, reference is made e.g. to the international application mentioned at the beginning, providing a more detailed description of collision situations in association with insertions, for example. Instead of conventional deletion updating, the structure may also employ functional updating implemented by known methods by copying the path from root to buckets.

As already stated at the beginning, the principles described above also relate to a bucketless trie structure. In such a case, the equivalent of a bucket is a data unit (to which a leaf node in the bucketless structure points). In the bucketless alternative, a problem will be presented by what type the root node must have when the structure only has one data unit (which is pointed to from the root node). In principle, there are two alternatives for implementing the root node. In accordance with the first alternative, the root may in this case have a quad node containing only one pointer. In accordance with the second, more advantageous alternative the root node is a compressed node directly referring to said data unit. Thus, in the case of one data unit the bucketless alternative must include this exception to the above rules.

Figure 11 shows a memory in accordance with the invention on block diagram level. Each dimension has a dedicated input register, and hence there is a total of n input registers. The search key of each dimension is stored in these input registers, denoted by references $R_1 \dots R_n$, each key in a register of its own. The input registers are connected to a register TR in which the above-described search word is formed in accordance with the bit interleaving method described above. The register TR is connected via adder S to the address input of memory MEM. The output of the memory in turn is connected to address register AR the output of which in turn is connected to adder S.

Initially the bits selected from each register are read into the common register TR in the correct order. In each ordinary quad node, two bits are selected from the registers, but with a compressed node the number of selected bits varies according to the number of bits to be searched in said node. The initial address of the first trie node is first stored in the address register AR, and the address obtained as an offset address from register TR is added to the initial address in adder S. The resulting address is supplied to the address input of the memory MEM, and the data output of the memory provides the initial address of the next trie node, the address being written into the address register AR over the previous address stored therein. Thereafter the next selected bits are again loaded from the input registers into the common register TR in the correct order, and the array address thus obtained is added to the initial address of the relevant array (i.e., trie node), obtained from the address register AR. This address is again supplied to the address input of the memory MEM, the data output of the memory thereafter providing the initial address of the next node. The above-described procedure is repeated until the desired point has been accessed and recordal can be performed or the desired record read.

Control logic CL attends to the maintenance of the memory structure and to the correct number of bits being extracted from the registers in each compressed node.

The rapidity of the address computation can be influenced by the type of hardware configuration chosen. Since progress is by way of the above-stated bit manipulations, address computation can be accelerated up by shifting from use of one processor to a multiprocessor environment in which parallel processing is carried out. An alternative implementation to the multiprocessor environment is an ASIC circuit.

Even though the invention has been described in the above with reference to examples in accordance with the accompanying drawings, it is obvious that the invention is not to be so restricted, but it can be modified within the scope of the inventive idea disclosed in the appended claims. The

5 address computation may also continue in the bucket, providing that unsearched bits remain. The structure may also be implemented for keys of variable length. The address computation may also continue in the bucket, providing that unsearched bits remain. The definition of a bucket given at the beginning is thus to be broadened to read that a bucket is a data structure that

10 may also contain another trie structure. Hence, several directory structures in accordance with the present invention can be linked in succession in such a way that another directory structure (that is, another trie structure) is stored in a bucket, or a pointer contained in a bucket or a leaf points to another directory structure. Reference from a bucket or a leaf is made directly to the root node of

15 the next directory structure. Generally, it may be stated that a bucket contains at least one element so that the type of an individual element is selected from a group comprising a data unit, a pointer to a stored data unit, a pointer to another directory structure and another directory structure. The detailed implementation of buckets is dependent on the application. In many cases, all

20 elements in buckets may be of the same type, being e.g. either a data unit or a pointer to a data unit. On the other hand, for instance in an application in which character strings are stored in the memory the bucket may contain element pairs in such a way that all pairs in the bucket are either pointer to data unit/pointer to directory structure pairs or data unit/pointer to a directory structure pairs or data unit/directory structure pairs. In such a case, for example, the

25 prefix of the character string may be stored in the data unit and the search may be continued from the directory structure that is the pair of the data unit.

Claims

1. A method for implementing a memory, in which memory data is stored as data units for each of which a dedicated storage space is assigned in the memory, in accordance with which method

5 - the memory is implemented as a directory structure comprising a tree-shaped hierarchy having nodes at several different levels, wherein an individual node can be (i) a trie node comprising an array wherein an individual element may contain the address of a lower node in the tree-shaped hierarchy and wherein an individual element may also be empty, or (ii) a bucket containing at least one element so that the type of an individual element in the

10 bucket is selected from a group including a data unit, a pointer to a stored data unit, a pointer to another directory structure and another directory structure,

 - address computation performed in the directory structure comprises the steps of

15 - (a) selecting in the node at the uppermost level of the tree-shaped hierarchy a predetermined number of bits from the bit string formed by the search keys employed, forming from the selected bits a search word with which the address of the next node is sought in the node, and proceeding to said node,

20 - (b) selecting from the unselected bits in the bit string formed by the search keys employed a predetermined number of bits and forming from the selected bits a search word with which the address of a further new node at a lower level is sought from the array of the node that has been accessed,

 - repeating step (b) until an empty element is encountered or until

25 the address of the new node at a lower level is the address of a bucket,

 c h a r a c t e r i z e d in that

 the arrays in the trie nodes are implemented as quad-nodes having a fixed size of four elements, and in at least part of the directory structure sets of successive quad nodes are replaced with compressed nodes in such a way

30 that an individual set made up by successive quad nodes, from each of which there is only one address to a quad node at a lower level, is replaced with a compressed node (CN) storing an address to the quad node that the lowest node in the set to be replaced points to, information on the value of the search word by means of which said address is found, and information on the total

35 number of bits from which search words are formed in the set to be replaced.

2. A method as claimed in claim 1, characterized in that replacement is carried out in the entire directory structure so that all said sets are replaced with compressed nodes.

5 3. A method as claimed in claim 1, characterized in that replacement is also carried out on a set having only one quad node, the total number of bits to be stored corresponding to the number of bits from which a search word is formed in said quad node.

10 4. A method as claimed in claim 1, characterized in that several successive compressed nodes are formed in the directory structure in such a way that at least in the compressed node at the uppermost level a number of search key bits to be searched corresponding to the word length employed is collected.

15 5. A method as claimed in claim 1, characterized in that several successive compressed nodes are combined into one new compressed node, the number of bits stored in the new node being the sum of the numbers obtained from the nodes to be combined.

20 6. A method as claimed in claim 4, characterized in that a chain made up by successive compressed nodes wherein the number of bits searched in at least two uppermost nodes corresponds to the word length employed is replaced with one collecting node (CN4) comprising

- an address to the node to which the lowest node in the chain contains an address,
- the sum of the numbers of searched bits obtained from the nodes in the chain, and
- 25 - the search word values contained in the chain nodes in sequence.

7. A method as claimed in claim 1, characterized in that in all uncompressed quad nodes of the memory, at least two addresses to a lower-level node are maintained.

30 8. A method for implementing a memory, in which memory data is stored as data units for each of which a dedicated storage space is assigned in the memory, in accordance with which method

- the memory is implemented as a directory structure comprising a tree-shaped hierarchy having nodes at several different levels, wherein an individual node can be (i) an internal node comprising an array wherein an individual element may contain the address of a lower node in the tree-shaped hierarchy and wherein an individual element may also be empty, or (ii) a leaf
- 35

containing at least one element the type of which is one from a group including a pointer to a stored data unit and a pointer to a node in another directory structure,

- address computation performed in the directory structure comprises the steps of
 - (a) selecting in the node at the uppermost level of the tree-shaped hierarchy a predetermined number of bits from the bit string formed by the search keys employed, forming from the selected bits a search word with which the address of the next node is sought in the node, and proceeding to said node,
 - (b) selecting from the unselected bits in the bit string formed by the search keys employed a predetermined number of bits and forming from the selected bits a search word with which the address of a further new node at a lower level is sought from the array of the node that has been accessed,
 - repeating step (b) until an empty element is encountered or until the address of the new node at a lower level is the address of a leaf, characterized in that the arrays in the internal nodes are implemented as quad nodes having a fixed size of four elements, and in at least part of the directory structure sets of successive quad nodes are replaced with compressed nodes in such a way that an individual set made up by successive quad nodes, from each of which there is only one address to a quad node at a lower level, is replaced with a compressed node (CN) storing an address to the quad node that the lowest node in the set to be replaced points to, information on the value of the search word by means of which said address is found, and information on the total number of bits from which search words are formed in the set to be replaced.

9. A method as claimed in claim 8, characterized in that replacement is performed in the entire directory structure in such a way that all said sets are replaced with compressed nodes.

10. A method as claimed in claim 8, characterized in that replacement is also carried out on a set comprising only one quad node, the total number of bits to be stored corresponding to the number of bits from which a search word is formed in said quad node.

11. A method as claimed in claim 8, characterized in that several successive compressed nodes are formed in the directory structure in

such a way that at least in the compressed node at the uppermost level a number of search key bits to be searched corresponding to the word length employed is collected.

- 5 12. A method as claimed in claim 8, characterized in that several successive compressed nodes are combined into one new compressed node, the number of bits stored in the new node being the sum of the numbers obtained from the nodes to be combined.

- 10 13. A method as claimed in claim 11, characterized in that a chain made up by successive compressed nodes wherein the number of bits searched in at least two uppermost nodes corresponds to the word length employed is replaced with one collecting node (CN4) comprising

- an address to the node to which the lowest node in the chain contains an address,
- the sum of the numbers of searched bits obtained from the nodes
- 15 in the chain, and
- the search word values contained in the chain nodes in sequence.

14. A method as claimed in claim 8, characterized in that in all uncompressed quad nodes of the memory, at least two addresses to a lower-level node are maintained.

1/5




FIG. 1





FIG. 2

2/5



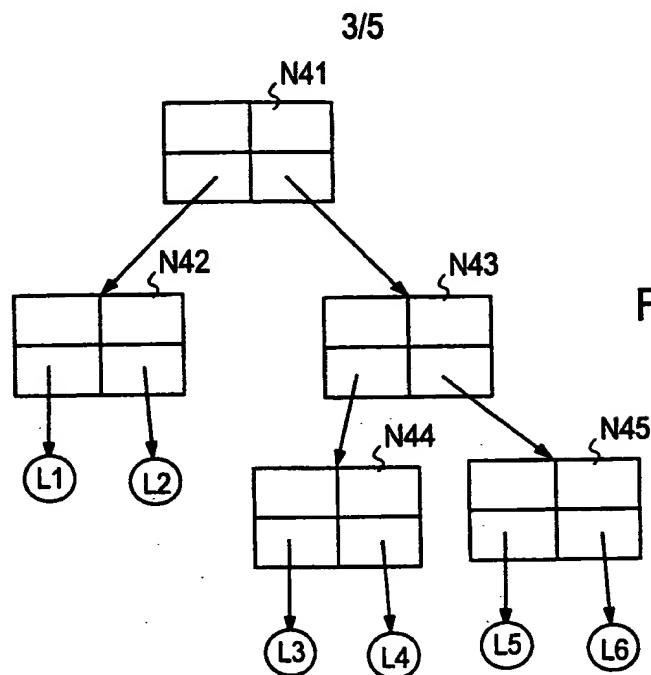


FIG. 4




FIG. 5




FIG. 6a




FIG. 6b

4/5




FIG. 7b

11	12	13	14	15
10	8	9	10	11
01	4	5	6	7
00	0	1	2	3
	00	01	10	11





FIG. 9

90	91	92
NUMBER OF SEARCHED BITS	ARRAY INDEX VALUE	POINTER TO CHILD NODE

5/5




FIG. 10b

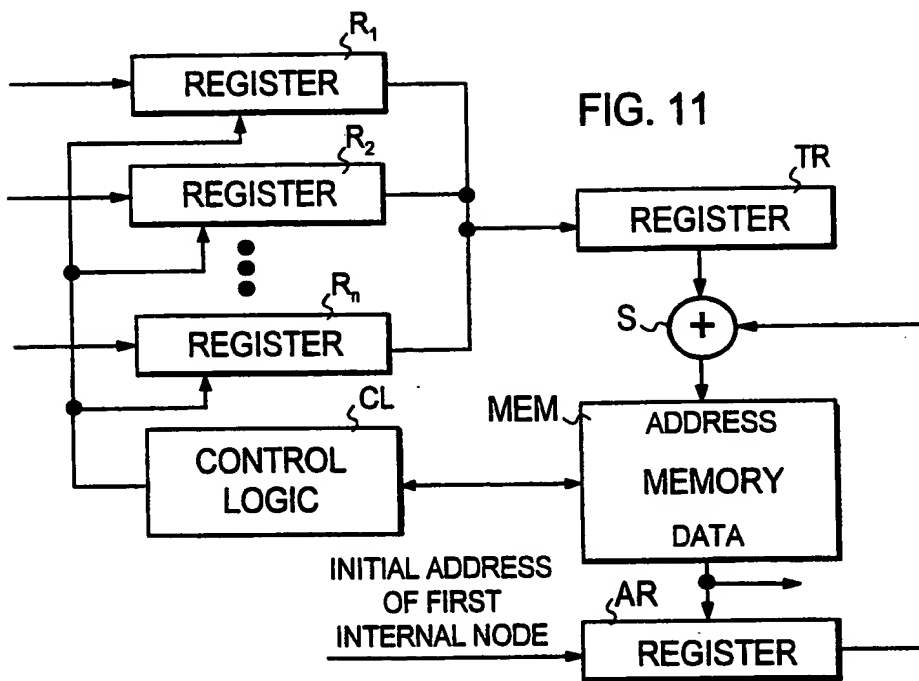


FIG. 11

INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI 98/00191

A. CLASSIFICATION OF SUBJECT MATTER

IPC6: G06F 17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC6: G06F, H04M

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WPI

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5276868 A (NIGEL T. POOLE), 4 January 1994 (04.01.94), column 1, line 13 - column 2, line 15; column 2, line 18 - column 3, line 30, figure 2 --	1-14
X	JP 8194719 A (FUJITSU LTD), 30 July 1996 (30.07.96), figures 11,12, abstract --	1-14
A	WO 9534155 A2 (NOKIA TELECOMMUNICATIONS OY), 14 December 1995 (14.12.95) --	1-14
A	WO 9600945 A1 (INTERNATIONAL BUSINESS MACHINES CORP.), 11 January 1996 (11.01.96) --	1-14

☒ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

18 August 1998

Date of mailing of the international search report

21 -08- 1998

Name and mailing address of the ISA/
Swedish Patent Office
Box 5055, S-102 42 STOCKHOLM
Facsimile No. +46 8 666 02 86

Authorized officer

Sylvain Dunand
Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.

PCT/FI 98/00191

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0650131 A1 (MICROSOFT CORPORATION), 26 April 1995 (26.04.95) -- -----	1-14

INTERNATIONAL SEARCH REPORT

Information on patent family members

27/07/98

International application No.

PCT/FI 98/00191

Patent document cited in search report			Publication date	Patent family member(s)		Publication date
US	5276868	A	04/01/94	CA	2043028 A	24/11/91
				EP	0458698 A	27/11/91
				JP	6004585 A	14/01/94

JP	8194719	A	30/07/96	US	5721899 A	24/02/98

WO	9534155	A2	14/12/95	AU	690282 B	23/04/98
				AU	2617495 A	04/01/96
				CN	1152365 A	18/06/97
				EP	0772836 A	14/05/97
				FI	942663 D	00/00/00
				JP	10504407 T	28/04/98
				FI	942664 D	00/00/00

WO	9600945	A1	11/01/96	EP	0804769 A	05/11/97

EP	0650131	A1	26/04/95	CA	2117846 A	21/04/95
				JP	7191891 A	28/07/95
				US	5752243 A	12/05/98

THIS PAGE BLANK (USPTO)